arXiv:1405.0707v1 [math.NA] 4 May 2014

PRECONDITIONING THE RESTARTED AND SHIFTED BLOCK FOM ALGORITHM FOR MATRIX EXPONENTIAL COMPUTATION

GANG WU*, HONG-KUI PANG†, AND JIANG-LI SUN ‡

Abstract. The approximation of $e^{tA}B$ where A is a large sparse matrix and B a rectangular matrix is the key ingredient in many scientific and engineering computations. A powerful tool to manage the matrix exponential function is to resort to a suitable rational approximation such as the Carathéodory-Fejér approximation, whose core reduces to solve shifted linear systems with multiple right-hand sides. The restarted and shifted block FOM algorithm is a commonly used technique for this problem. However, determining good preconditioners for shifted systems that preserve the original structure is a difficult task. In this paper, we propose a new preconditioner for the restarted and shifted block FOM algorithm. The key is that the absolute values of the poles of the Carathéodory-Fejér approximation are medium sized and can be much smaller than the norm of the matrix in question. The advantages of the proposed strategy are that we can precondition all the shifted linear systems simultaneously, and preserve the original structure of the shifted linear systems after restarting. Theoretical results are provided to show the rationality of our preconditioning strategy. Applications of the new approach to Toeplitz matrix exponential problem are also discussed. Numerical experiments illustrate the superiority of the new algorithm over many state-of-the-art algorithms for matrix exponential.

Key words. Matrix exponential, Shifted linear systems, Multiple right-hand sides, Full Orthogonalization Method (FOM), Precondition, Carathéodory-Fejér approximation.

AMS subject classifications. 65F60, 65F10, 65F15.

1. Introduction.

The computation of exponential function of large and sparse matrices is now ubiquitous in numerical analysis research [33, 43]. In this paper, we pay special attention to numerical approximation of the matrix exponential to a block vector

$$Z(t) = e^{tA}B, \tag{1.1}$$

where $t \in \mathbb{R}$ is a fixed constant, $tA \in \mathbb{R}^{n\times n}$ is an n-by-n large sparse, negative definite matrix with $\|tA\| \gg 1$, and $B = \left[\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_p\right] \in \mathbb{R}^{n\times p}$ is a block vector with $1 \leq p \ll n$. That is, the real part of the spectrum of tA lies in the left half plane. This problem is the key ingredient of the computation of exponential integrators [3, 13, 35], time-dependent partial differential equations [26], the approximation of dynamical systems [4], and so on [5, 38, 66]. For notation simplicity, we denote tA by A from now on unless otherwise stated.

The Krylov subspace methods are widely used for this type of problem, which work on the subspaces generated by A and $\mathbf{b}_i$ $(i = 1, 2, \ldots, p)$. Generally speaking, there are two classes of Krylov subspace methods for evaluating (1.1) when A is large

*Corresponding author (G. Wu). Department of Mathematics, China University of Mining and Technology, Xuzhou, 221116, P. R. China. E-mail: wugangzy@gmail.com and gangwu76@126.com. This author is supported by the National Science Foundation of China under grant 11371176, the Natural Science Foundation of Jiangsu Province under grant BK20131126, the 333 Project of Jiangsu Province, and the Talent Introduction Program of China University of Mining and Technology.

†School of Mathematics and Statistics, Jiangsu Normal University, Xuzhou, 221116, Jiangsu, P. R. China. E-mail: panghongkui@163.com. This author is supported by the National Science Foundation of China under grant 11201192, the Natural Science Foundation of Jiangsu Province under grant BK2012577, and the Natural Science Foundation for Colleges and Universities in Jiangsu Province under grant 12KJB110004.

‡School of Mathematics and Statistics, Jiangsu Normal University, Xuzhou, 221116, Jiangsu, P. R. China. E-mail: sunjiangli1108@163.com.

and sparse [50]. In the first class of methods, the matrix is projected into a much smaller subspace, then the exponential is applied to the reduced matrix, and finally the approximation is projected back to the original large space [10, 17, 18, 19, 44, 53, 70, 75]. In the second class of methods, the exponential function is first approximated by simpler functions such as rational functions, and then the action of the matrix exponential is evaluated [23, 24, 37, 67, 68, 74]. In this case, the core of (1.1) reduces to solving the following linear systems with multiple shifts and multiple right-hand sides

$$(A - \tau_i I)X_i = B, \quad i = 1, 2, \ldots, \nu, \tag{1.2}$$

where $\tau_i \in \mathbb{C}$ $(i = 1, 2, \ldots, \nu)$ are the shifts. Throughout this paper, we make the assumption that both A and the shifted matrices $A - \tau_i I$ $(i = 1, 2, \ldots, \nu)$ are nonsingular. Note that there are $\nu \times p$ linear systems altogether.

The use of rational functions and their partial fraction expansion allows us to exploit and generalize known properties of Krylov subspace methods for the solution of algebraic linear systems [52, 60]. For instance, the first way for solving (1.2) is to evaluate the p shifted linear systems separately using some shifted Krylov subspace algorithms [21, 22, 30, 58, 63]. In [16], Darnnu *et al.* proposed a deflated and shifted GMRES algorithm for systems with multiple shifts and multiple right-hand sides. In this algorithm, the p shifted linear systems are solved separately, and eigenvector information from solution of the first right-hand side is utilized to assist the convergence of the subsequent ones. However, one has to solve the shifted systems with an auxiliary right-hand side, which is unfavorable for very large matrices. The second way for solving (1.2) is to calculate the ν linear systems with multiple right-hand sides sequentially, using some block Krylov subspace methods such as the block GMRES algorithm with deflation [46]. However, one has to solve the ν linear systems with multiple right-hand sides one by one, and this algorithm may suffer from the drawback of slow convergence [46].

The third way is to use the shifted block GMRES algorithm [73]. In essence, it is a generalization of the shifted GMRES algorithm [22] to its block version. There are some deficiencies in the shifted block GMRES algorithm for linear systems with complex shifts. Firstly, when both A and B are real while the shifts τ_i $(i = 1, 2, \ldots, \nu)$ are complex, the initial block vector for the next cycle after restarting is also complex. This problem can arise in projection methods for rational approximation to the matrix exponential [37, 74]. Consequently, the expensive steps for constructing the orthogonal basis have to be performed in complex arithmetic after the first cycle. Secondly, the shifted block GMRES algorithm may suffer from the drawbacks of "stagnation" and "near breakdown" in practice [73]. Thirdly, one has to pick the "seed" block system in advance. More precisely, we have to determine τ_1 in advance, which is a difficult task if there is no further knowledge on the shifted block linear systems *a priori*. Finally, in the shifted block GMRES algorithm, only the "seed" block system residual is minimized, whereas no minimization property is satisfied by the residuals of the "additional" block systems [22, 73].

In order to overcome these difficulties, we first introduce a shifted block Full Orthogonalization Method (FOM) with deflated restarting. It is a block and deflated version of the shifted FOM algorithm due to Simoncini [58]. However, the algorithm may converge very slowly or even may fail to converge in practice. A commonly used technique for accelerating iterative solvers is preconditioning [52, 60]. In general, the reliability of iterative techniques, when dealing with various applications, depends

much more on the quality of the preconditioner than on the particular Krylov subspace methods [52]. Therefore, how to construct an efficient preconditioner is a crucial task for shifted linear systems with multiple right-hand sides [16, 31, 60].

The study of preconditioners for shifted linear systems or sequences of linear systems is an active area, and many technologies have been proposed when A is symmetric positive definite (SPD), complex symmetric or Hermitian; see [6, 7, 8, 25, 64] and the references therein. However, when A is non-Hermitian, there are two major limitations to preconditioning the shifted linear systems, which diminish the usefulness of precondition techniques considerably [31]. First, we have to start with the same residual for all values of τ_i, which means that we cannot have τ_i-dependent left preconditioning. Second, preconditioning must retain the shifted structure of the matrix. Indeed, as was pointed out by Simoncini and Szyld [60], standard preconditioning approaches may be effective on each shifted system; however, they destroy the shifted structure so that the convenient invariance property of the Krylov subspace can no longer be employed. In this case, determining good preconditioners for shifted systems that preserve the original structure is very important and is still an open area of research [60].

In [30], Gu *et al.* presented a flexible preconditioner for solving the shifted systems. This method allows one to incorporate different preconditioners $(A - \tau_i I)^{-1}$ with different τ_i, into the Arnoldi procedure when constructing a projective subspace. Unfortunately, one has to solve m (complex) shifted systems $(A - \tau_i I)\mathbf{w}_i = \mathbf{v}_i$ $(i = 1, 2 \ldots, m)$ in each outer iteration of the algorithm, where m is the number of Arnoldi steps. Thus, the expensive step of constructing the orthogonal basis can not be realized in real arithmetics even if A and B are real. In particular, when direct methods are used, an additional difficulty is that for solving a sequence of (complex) shifted systems, the matrix has to be refactorized for each shift, resulting in considerable computational effort. Just recently, this strategy was applied to precondition shifted linear systems arising in oscillatory hydraulic tomography (OHT) [54]. In [2, 31], some polynomial preconditioners are investigated to solve shifted linear systems together with the bi-conjugate gradient method (BiCG). Unfortunately, for a general polynomial preconditioner, systems with different shifts may no longer have equivalent Krylov subspaces for the shifted FOM (or GMRES) algorithm. Furthermore, when the eigenvalues of A are complex, the foci of the ellipse that encloses the spectrum has to be determined in advance [2]. Recently, a preconditioned and shifted (block) GMRES algorithm was proposed for accelerating the PageRank computation [73]. In this approach, a polynomial preconditioner was introduced for shifted linear systems with multiple right-hand sides, under the condition that the spectrum radius $\rho(A) < |\tau_i|,\ i = 1, 2, \ldots, \nu$. In many applications, however, we have that $\|A\| \gg |\tau_i|$ [4, 33, 36, 49, 66, 69]. In this situation, the preconditioning strategy proposed in [73] does not work any more.

In this work, we propose a new preconditioner for shifted linear systems with multiple right-hand sides when $\|A\| \gg |\tau_i|,\ i = 1, 2, \ldots, \nu$, and present a preconditioned and shifted block FOM algorithm with deflated restarting for the matrix exponential problem. The motivation of our new method is threefold: First, we aim to precondition all the shifted linear system *simultaneously*. Second, we are interested in solving all the shifted linear systems in the same search subspace. Third, when both A and B are real while the shifts are complex, the expensive step of constructing the orthogonal basis can be realized in real arithmetics.

This paper is organized as follows. In Section 2, we briefly introduce the Caratheé

odory-Fejér method for rational approximation of exponential [56, 67, 68, 71]. In Section 3, we present a shifted block FOM algorithm with deflated restarting. In Section 4, we propose a new preconditioner and provide a preconditioned and shifted block FOM algorithm with deflation for the matrix exponential problem (1.2). Theoretical analysis is given to show the feasibility and rationality of the new approach. Numerical experiments are reported in Section 5, illustrating that our new algorithm often outperforms many state-of-the-art algorithms for matrix exponential. Some concluding remarks are given in Section 6.

In this paper, matrices are written as capitals such as A, and vectors are typeset in bold such as $\mathbf{v}$. Let $\mathcal{K}_m(A, V_1) = \text{span}\{V_1, AV_1, \ldots, A^{m-1}V_1\}$ be the block Krylov subspace generated by A and V_1, let A^{T} and A^{H} be the transpose and conjugate transpose of A, respectively. Denote by $\|\cdot\|$ the vector norm or its induced matrix norm, and by $\|\cdot\|_F$ the F-norm of a matrix, respectively. Let I be the identity matrix and O be a zero matrix or vector whose order are clear from the context. Let E_j be an $mp \times p$ matix which consists of the $\big((j-1)p+1\big)$-th to the jp-th $(1 \le j \le m)$ columns of the $mp \times mp$ identity matrix. Standard MATLAB [40] notations are used whenever necessary.

2. The Carathéodory-Fejér method for rational approximation. A very powerful tool to manage matrix functions is to resort to a suitable rational approximation [33, 37, 67, 68, 71, 74]. Given a function $f(x)$ which is continuous on a closed interval, let $\mathcal{R}_{\mu,\nu}$ be the set of rational functions of the form

$$\widehat{f}(x) = \frac{\sum_{k=0}^{\mu} a_k x^k}{\sum_{k=0}^{\nu} b_k x^k}. \tag{2.1}$$

The well known Chebyshev approximation consists in finding an approximation $f^\star(x) \in \mathcal{R}_{\mu,\nu}$, such that [27, 33, 42]

$$\|f(x) - f^\star(x)\|_\infty \le \|f(x) - \widehat{f}(x)\|_\infty, \quad \forall \widehat{f}(x) \in \mathcal{R}_{\mu,\nu}, \tag{2.2}$$

where $\|f(x) - f^\star(x)\|_\infty$ stands for the infinity norm of $f(x) - f^\star(x)$. If $\mu = \nu$ and $f^\star(x)$ is the Chebyshev rational approximation, it holds that [12]

$$\sup_{x \ge 0} \left|e^{-x} - f^\star(x)\right| \approx 10^{-\nu}.$$

It is shown that $f^\star(x)$ always exists [42], however, its uniqueness is not guaranteed unless $\nu = 0$. Furthermore, how to characterize the best approximation remains an issue [37, 42].

A practical alternative is to use the Carathéodory-Fejér (CF) approximation, and to work directly with rational functions designed as good approximations to e^x on $\mathbb{R}^-$ (i.e., the left half plane) [56, 67, 68, 71]. This technique is based on the singular value decomposition of a Hankel matrix of Chebyshev coefficients of the function e^x, transplanted from $\mathbb{R}^-$ to $[-1, 1]$. The goal of this method is to find a function $\widetilde{f}^\star(x) \in \widetilde{\mathcal{R}}_{\mu,\nu}$ such that

$$\|f(x) - \widetilde{f}^\star(x)\|_\infty \le \|f(x) - \widetilde{f}(x)\|_\infty, \quad \forall \widetilde{f}(x) \in \widetilde{\mathcal{R}}_{\mu,\nu}, \tag{2.3}$$

where $\widetilde{\mathcal{R}}_{\mu,\nu}$ is the set of rational functions of the following form [27, 67, 68, 71]

$$\widetilde{f}(x) = \frac{\sum_{k=-\infty}^{\mu} a_k x^k}{\sum_{k=0}^{\nu} b_k x^k}. \tag{2.4}$$

Then the Carathéodory-Fejér approximation is obtained from dropping all terms of negative degree in the numerator of $\widetilde{f}^{\star}(x)$.

Although the Carathéodory-Fejér approximation is in principle only approximate, Magnus [39] indicated that for $\mu = \nu$, the CF approximation differs from the best Chebyshev approximation for a factor of $\mathcal{O}(56^{-\nu})$. Indeed, the error in exact arithmetic is below standard machine precision for $\nu \geq 9$, and for $\nu = 14$ it is about 10^{-26} [39, 68].

Without loss of generality, we still denote by

$$\widetilde{f}^{\star}(x) = \omega_0 + \sum_{i=1}^{\nu} \frac{\omega_i}{x - \tau_i}$$

the Carathéodory-Fejér approximation for approximating e^x on $\mathbb{R}^-$, where τ_i and ω_i $(i = 1, 2, \ldots, \nu)$ are the poles and the residues, respectively. We refer to [68] for a MATLAB function to compute the poles and the residues of the Carathéodory-Fejér method for the type (ν, ν) near best approximation. The idea behind the Carathéodory-Fejér approach for matrix exponential is to approximate $e^A B$ by using

$$\widetilde{f}^{\star}(A)B = \omega_0 B + \sum_{i=1}^{\nu} \omega_i (A - \tau_i I)^{-1} B. \tag{2.5}$$

Thus, an approximation to $\widetilde{f}^{\star}(A)B$ can be obtained by first solving the shifted linear systems appearing in the sum, and then by collecting the sum terms. When the shifts $\{\tau_i\}_{i=1}^{\nu}$ are appeared in conjugate pairs, one only needs to solve $\nu/2$ shifted linear system altogether.

REMARK 2.1. *In the Carathéodory-Fejér approximation of the matrix exponential, we point out that the absolute values of the shifts τ_i $(i = 1, 2, \ldots, \nu)$ are medium sized. For instance, if we choose $\nu = 14$ in the Carathéodory-Fejér approximation, then $\max_{1 \leq i \leq \nu} |\tau_i| \approx 18.9$. In many applications, we have that $\|A\| \gg |\tau_i|$ $(i = 1, 2, \ldots, \nu)$ [4, 33, 36, 49, 66, 69].*

3. A shifted block FOM algorithm with deflated restarting for shifted linear systems with multiple right-hand sides. A popular way to solve linear system with multiple right-hand sides is the block approach [52, 60]. In this section, we introduce a shifted block FOM algorithm with the deflated restarting strategy [46]. It can be viewed as a block and delated version of the shifted FOM algorithm given by Simoncini [58].

Suppose that V_1 is an $n \times p$ orthonormal matrix, then the block Arnoldi process with A. Ruhe's variant [52] will generate an orthnormal basis for the block Krylov subspace $\mathcal{K}_{m+1}(A, V_1) = \text{span}\{V_1, AV_1, \ldots, A^m V_1\}$ in exact arithmetics. The following relation holds for the block Arnoldi process [52]

$$A\mathcal{V}_m = \mathcal{V}_m \mathcal{H}_m + V_{m+1} H_{m+1,m} E_m^{\mathrm{T}} = \mathcal{V}_{m+1} \widetilde{\mathcal{H}}_m, \tag{3.1}$$

where $\mathcal{V}_{m+1} = \big[V_1, V_2, \ldots, V_{m+1}\big]$ is an orthogonal basis for $\mathcal{K}_{m+1}(A, V_1)$, E_m is an $mp \times p$ matix which consists of the last p columns of the $mp \times mp$ identity matrix, and $\mathcal{H}_m, \widetilde{\mathcal{H}}_m$ are $mp \times mp$ and $(m+1)p \times mp$ upper band Hessenberg matrices, respectively, with the $p \times p$ matrices $H_{i,j}$ being their elements.

It is known that the block Krylov subspace is spanned by the same basis if the matrix is scaled or shifted [52]

$$\text{span}\{V_1, (A - \tau_i I)V_1, \ldots, (A - \tau_i I)^m V_1\} = \text{span}\{V_1, AV_1, \ldots, A^m V_1\},\ i = 1, 2, \ldots, \nu.$$

By (3.1), the block Arnoldi relations with respect to the shifted matrices $A-\tau_i I$ are

$$\begin{aligned}(A-\tau_i I)\mathcal{V}_m &= \mathcal{V}_m(\mathcal{H}_m-\tau_i I)+V_{m+1}H_{m+1,m}E_m^{\mathrm{T}}\\ &= \mathcal{V}_{m+1}(\widetilde{\mathcal{H}}_m-\tau_i\widetilde{I}_m),\quad i=1,2,\ldots,\nu,\end{aligned}\tag{3.2}$$

where $\widetilde{I}_m$ is an $(mp+p)\times mp$ matrix which is the same as the identity matrix except for p zero rows at the bottom. However, the shifted block FOM algorithm will become impractical when m is large, because of the growth of memory requirement and computation cost.

One remedy is to restart the shifted block FOM algorithm. Denote by $\widehat{X}_i^0$ the approximations obtained from the "previous" cycle (where we choose $\widehat{X}_i^0=O$ for the first cycle), and by $\widehat{R}_i^0=B-(A-\tau_i I)\widehat{X}_i^0$ the corresponding residuals. Then the shifted block FOM algorithm uses $\widehat{X}_i^{FOM}=\widehat{X}_i^0+\mathcal{V}_m\widehat{Z}_i^{FOM}$ $(i=1,2,\ldots,\nu)$ as approximate solutions to (1.2) in the "current" cycle, with the residuals

$$\widehat{R}_i^{FOM}=B-(A-\tau_i I)\widehat{X}_i^{FOM}\perp \mathrm{span}\{\mathcal{V}_m\},\quad i=1,2,\ldots,\nu.\tag{3.3}$$

Let $\widehat{R}_i^0=V_1\widehat{R}_i$ be the QR factorizations of $\widehat{R}_i^0$, then it follows from (3.3) that $\widehat{Z}_i^{FOM}$ can be obtained from solving the following *projected* linear systems

$$(\mathcal{H}_m-\tau_i I)\widehat{Z}_i^{FOM}=E_1\widehat{R}_i,\quad i=1,2,\ldots,\nu,\tag{3.4}$$

where E_1 is the matrix composed of the first p columns of the $mp\times mp$ identity matrix. Note that the residuals with respect to different shifts are included in $\mathrm{span}\{V_{m+1}\}$ in the shifted block FOM algorithm. Indeed, we have from (3.2) that

$$\begin{aligned}\widehat{R}_i^{FOM} &= B-(A-\tau_i I)\widehat{X}_i^{FOM}=\widehat{R}_i^0-(A-\tau_i I)\mathcal{V}_m\widehat{Z}_i^{FOM}\\ &= \mathcal{V}_m\big[E_1\widehat{R}_i-(\mathcal{H}_m-\tau_i I)\widehat{Z}_i^{FOM}\big]-V_{m+1}(H_{m+1,m}E_m^{\mathrm{T}}\widehat{Z}_i^{FOM})\\ &= -V_{m+1}(H_{m+1,m}E_m^{\mathrm{T}}\widehat{Z}_i^{FOM}),\quad i=1,2,\ldots,\nu,\end{aligned}\tag{3.5}$$

and

$$\|\widehat{R}_i^{FOM}\|_F=\|H_{m+1,m}E_m^{\mathrm{T}}\widehat{Z}_i^{FOM}\|_F,\quad i=1,2,\ldots,\nu.\tag{3.6}$$

As a result, the residuals with respect to different shifts are collinear in the shifted block FOM algorithm.

Now we consider how to restart the shifted block FOM algorithm in practice. In this paper, we are interested in the deflation strategy [46] for linear systems, in which the approximate eigenvectors are put firstly in the search subspace. Here "deflation" means computing eigenvectors corresponding to some eigenvalues, and using them to remove these eigenvalues from the spectrum of the matrix, to speed up the convergence of the iterative algorithm. More precisely, let k be a multiple of p and let $\widetilde{\mathbf{y}}_1,\widetilde{\mathbf{y}}_2,\ldots,\widetilde{\mathbf{y}}_k$ be k Ritz vectors [29] computed from the "previous" cycle. Then after restarting, the search subspace of the "current" cycle is [46]

$$\mathrm{span}\{\widetilde{\mathbf{y}}_1,\widetilde{\mathbf{y}}_2,\ldots,\widetilde{\mathbf{y}}_k,V_{m+1},AV_{m+1},\ldots,A^qV_{m+1}\},\tag{3.7}$$

where $q=(mp-k)/p$. The shifted block FOM algorithm with deflation generates a block Arnoldi relation similar to (3.1), where $\mathcal{V}_m$ is an $n\times mp$ matrix whose columns span the subspace (3.7), and $\widetilde{\mathcal{H}}_m$ is an $(mp+p)\times mp$ matrix that is band upper-Hessenberg except for a full $(k+p)\times k$ leading portion $\widetilde{\mathcal{H}}_k^{new}$. A part of this recurrence

can be separated out to give $A\mathcal{V}_k^{new} = \mathcal{V}_{k+p}^{new}\widetilde{\mathcal{H}}_k^{new}$, where $\mathcal{V}_{k+p}^{new}$ is an $n\times(k+p)$ matrix whose columns span the subspace of Ritz vectors and V_{m+1}, and $\mathcal{V}_k^{new}$ consists of the first k columns of $\mathcal{V}_{k+p}^{new}$; for more details, refer to [32, 46]. Applying this deflation strategy to the shifted block FOM algorithm introduced above, we can present the following algorithm.

ALGORITHM 1. A shifted block FOM algorithm with deflation for shifted linear systems with multiple right-hand sides **(SBFOM-DR)**

1. *Input: Given the block Arnoldi steps m, a prescribed tolerance tol, and k (which is a multiple of p), the number of approximate eigenvectors retained from the previous cycle, and set all the initial guess $\widehat{X}_i^0 = O,\ i = 1, 2, \ldots, \nu$;*

2. *Compute the QR decomposition: $\widehat{R}_1^0 = V_1\widehat{R}_1$;*

3. *Use V_1 as the initial block vector, and run the block Arnoldi process (with A. Ruhe's variant) for the computation of $\mathcal{V}_m$ and $\widetilde{\mathcal{H}}_m$;*

4. *Solve the projected shifted linear systems. If all the residual norms are below the prescribed tolerance tol, see (3.6), then Stop, else Continue;*

5. *Compute all the eigenpairs $(\widehat{\theta}_i, \mathbf{y}_i)$ $(i = 1, 2, \ldots, mp)$ of $\mathcal{H}_m$, and select k smallest of them as the desired ones;*

6. *Orthonormalization of the first k short vectors: Orthonormalize the $\{\mathbf{y}_i\}$'s, first separating them into real parts and imaginary parts if they are complex, in order to form an $mp \times k$ matrix $P_k = [\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_k]$. Both parts of complex vectors need to be included, so k may be increased or decreased if necessary;*

7. *Orthonormalization of the $k + p$ short vectors: Extend P_k to an $(mp + p) \times k$ matrix $\widehat{P}_k$, by appending a $p \times k$ zero matrix at the bottom, and set $P_{k+p} = [\widehat{P}_k,\ \widehat{E}]$, where $\widehat{E}$ is an $(mp + p) \times p$ zeros matrix except for its last p rows being an identity matrix. Note that P_{k+p} is an $(mp + p) \times (k + p)$ matrix;*

8. *Form portions of new $\widetilde{\mathcal{H}}_{mp}$ and $\mathcal{V}_m$ using the old $\widetilde{\mathcal{H}}_{mp}$ and $\mathcal{V}_{m+1}$: Let $\widetilde{\mathcal{H}}_k^{new} = P_{k+p}^{\mathrm{T}}\widetilde{\mathcal{H}}_{mp}P_k$, and $\mathcal{V}_{k+p}^{new} = \mathcal{V}_{m+1}P_{k+p}$;*

9. *Block Arnoldi iteration: Apply the block Arnoldi iteration from the current point to form the rest portions of $\mathcal{V}_{m+1}$ and $\widetilde{\mathcal{H}}_{mp}$, where the current point is $\mathcal{V}_{k+p}^{new}$, and goto* **Step 4**.

Algorithm 1 is very attractive when both A and B are real while the shifts $\{\tau_i\}$'s are complex. Indeed, at each cycle after restarting, all the complex residuals are collinear to the $(m + 1)$-st real basis vector V_{m+1}, and the expensive step for constructing the orthogonal basis $\mathcal{V}_{m+1}$ can be performed in real arithmetics.

4. Preconditioning the restarted and shifted block FOM algorithm.

In this section, we propose a new preconditioner for shifted linear systems with multiple right-hand sides when $\|A\| \gg |\tau_i|$ $(i = 1, 2, \ldots, \nu)$, and apply it to the computation of the matrix exponential problem. Some theoretical results are established to show the rationality and feasibility of our new preconditioning strategy.

4.1. The preconditioner and the algorithm.

In many applications, we have that $\|A\| \gg \max_{1\le i\le\nu} |\tau_i|$ [4, 33, 36, 49, 66, 69]. For instance, if A is the usual finite difference/element discretization of a Poisson operator with Dirichlet boundary conditions, the norm of A grows rapidly with respect to the number n of spatial grid points [51]. Indeed, if $\|A\| > |\tau_i|$, then

$$\frac{\|(A - \tau_i I) - A\|}{\|A - \tau_i I\|} \le \frac{|\tau_i|}{\|A\| - |\tau_i|}, \quad i = 1, 2, \ldots, \nu.$$

Specifically, if $\|A\| \gg |\tau_i|$, we have that

$$\frac{\|(A-\tau_i I)-A\|}{\|A-\tau_i I\|} \lessapprox \frac{|\tau_i|}{\|A\|} \ll 1, \quad i=1,2,\ldots,\nu. \tag{4.1}$$

Thus, the idea is to use

$$\widetilde{M}_i = A^{-1}, \quad i=1,2,\ldots,\nu, \tag{4.2}$$

as right-preconditioner to (1.2). In other words, the matrix A can be viewed as an approximation to $A-\tau_i I$ if $\|A\| \gg |\tau_i|$ $(i=1,2,\ldots,\nu)$, and our strategy can be understood as an approximate inverse preconditioning or a shift-invert technique where the shift is set to be zero. Recall from Remark 2.1 that the assumption of $\|A\| \gg \max_{1\le i\le\nu}|\tau_i|$ is mild for the Carathéodory-Fejér approximation of the matrix exponential. Moreover, it is seen that

$$\|I-(A-\tau_i I)A^{-1}\| = \|\tau_i A^{-1}\| = \frac{|\tau_i|}{\|A\|}\cdot\kappa(A), \quad i=1,2,\ldots,\nu, \tag{4.3}$$

where $\kappa(A) = \|A\|\cdot\|A^{-1}\|$ is the condition number of A. Therefore, A^{-1} will be a good preconditioner to the shifted linear systems if $\max_{1\le i\le\nu}\frac{|\tau_i|}{\|A\|}$ is sufficiently small and A is not too ill-conditioned. Furthermore, the larger $\|A\|$ is, the better the preconditioner will be. Indeed, using A^{-1} as the preconditioner may be advantageous with respect to $(A-\tau_i I)^{-1}$ if either of the following conditions apply [59]: First, solving with $A-\tau_i I$ is more expensive or difficult than solving with A (e.g., A is large and real while the $\{\tau_i\}$'s are complex). Second, the number of shifts involved is large, and the dimension of the search space is significantly low. Both conditions can be satisfied in many real applications [18, 59].

Therefore, the preconditioned linear systems turn out to be the following form

$$(A-\tau_i I)\widetilde{M}_i Y_i = B, \quad i=1,2,\ldots,\nu,$$

that is,

$$\big(I-\tau_i A^{-1}\big)Y_i = B, \quad i=1,2,\ldots,\nu, \tag{4.4}$$

with $X_i = \widetilde{M}_i Y_i = A^{-1}Y_i$ $(i=1,2,\ldots,\nu)$ being the desired solutions.

REMARK 4.1. *It is known that clustered eigenvalues are favorable for convergence of Krylov subspace methods [52]. Let λ_j be eigenvalues of A, then the eigenvalues of $(I-\tau_i A^{-1})$ are mapped to be $1-\tau_i/\lambda_j$. If the smallest eigenvalues of A are deflated using the technique described in Section 3, then the eigenvalues of $(I-\tau_i A^{-1})$ will be clustered around 1 as $|\lambda_j| \gg |\tau_i|$. This gives the rationality of knitting the deflated restarting technique together with our preconditioning strategy.*

We notice that the block Krylov subspace with respect to A^{-1} and those with respect to $(A-\tau_i I)\widetilde{M}_i$ $(i=1,2,\ldots,\nu)$ are identical. Moreover, suppose that the block Arnoldi relation with respect to A^{-1} is

$$A^{-1}\mathcal{V}_m = \mathcal{V}_m\mathcal{H}_m + V_{m+1}H_{m+1,m}E_m^{\mathrm{T}} = \mathcal{V}_{m+1}\widetilde{\mathcal{H}}_m. \tag{4.5}$$

Recall that both $\mathcal{V}_{m+1}$ and $\widetilde{\mathcal{H}}_m$ are different from those in (3.1). Then, the block Arnoldi relations with respect to the shift systems (4.4) are

$$\big(I-\tau_i A^{-1}\big)\mathcal{V}_m = \mathcal{V}_{m+1}\big(\widetilde{I}_m - \tau_i\widetilde{\mathcal{H}}_m\big), \quad i=1,2,\ldots,\nu. \tag{4.6}$$

So we can precondition all the shifted linear systems simultaneously, and solve them in the same search subspace in the preconditioned and shifted block FOM algorithm.

Denote by $\widehat{Y}_i^0$ the approximations obtained from the "previous" cycle of the preconditioned and shifted block FOM algorithm, and by $\widehat{R}_i^0 = B - (I - \tau_i A^{-1})\widehat{Y}_i^0$ $(i = 1, 2, \ldots, \nu)$ the corresponding residuals. Let $\widehat{R}_i^0 = V_1 \widehat{R}_i$ be the QR factorization, then the preconditioned and shifted block FOM algorithm uses $\widehat{Y}_i^{FOM} = \widehat{Y}_i^0 + \mathcal{V}_m \widehat{Z}_i^{FOM}$ as approximate solutions to (4.4), where $\widehat{Z}_i^{FOM}$ are the solutions of the following *projected* linear systems

$$(I - \tau_i \mathcal{H}_m)\widehat{Z}_i^{FOM} = E_1 \widehat{R}_i, \quad i = 1, 2, \ldots, \nu. \tag{4.7}$$

The residuals are

$$\begin{aligned} \widehat{R}_i^{FOM} &= B - (I - \tau_i A^{-1})(\widehat{Y}_i^0 + \mathcal{V}_m \widehat{Z}_i^{FOM}) \\ &= \tau_i V_{m+1}(H_{m+1,m} E_m^{\mathrm{T}} \widehat{Z}_i^{FOM}), \quad i = 1, 2, \ldots, \nu, \end{aligned} \tag{4.8}$$

and

$$\|\widehat{R}_i^{FOM}\|_F = |\tau_i| \cdot \|H_{m+1,m} E_m^{\mathrm{T}} \widehat{Z}_i^{FOM}\|_F, \quad i = 1, 2, \ldots, \nu. \tag{4.9}$$

We are ready to propose the main algorithm of this paper for shifted linear systems with multiple right-hand sides as $\|A\| \gg |\tau_i|,\ i = 1, 2, \ldots, \nu$.

ALGORITHM 2. A preconditioned and shifted block FOM algorithm with deflation for shifted linear systems with multiple right-hand sides **(PSBFOM-DR)**
Steps **1**–**2** *are the same as those in Algorithm 1;*
3. *Run the block Arnoldi process* (*with A. Ruhe's variant*) *for the computation of* (*4.6*)*;*
4. *Solving the projected shifted linear systems. If the residual norms are below the prescribed tolerance tol, see* (4.9*), then Stop and form the numerical solutions* $\widetilde{X}_i$ $(i = 1, 2, \ldots, \nu)$*, else continue;*
5. *Steps* **5**–**9** *are similar to those in Algorithm 1, except for the block Krylov subspace is generated by using* A^{-1}*.*

In each cycle, some linear systems with respect to A must be solved in the block Arnoldi process to expand the basis. More precisely, mp for the first cycle and $mp - k$ for the cycles after the first. This makes the generation of the search space more expensive than the unpreconditioned one. In practical calculations, if A is of medium sized (say, $n \leq 10,000$), we can perform the sparse LU factorization of A for the inverse. Moreover, the LU factorization requires to be performed once and for all, and the L and U factors can be stored for later use. Specifically, when A has some special structure such as Toeplitz, one can use the Gohberg–Semencul formula [28] for the matrix-vector products with respect to A^{-1}, with no need to form and store the inverse explicitly; see [36, 48] for more details.

On the other hand, if the matrix is so large that direct methods are prohibitive, using an iterative solver is definitely possible and advisable [52, 60]. Another alternative is to use some inexact Krylov subspace algorithms or the flexible preconditioning strategy in which the products with respect to A^{-1} can be computed inaccurately as the outer iteration converges [11, 61]. In [54], an inexact preconditioning was considered for solving the shifted linear systems. One can naturally extend the idea of inexact preconditioning to systems with multiple shifts and multiple right-hand sides using block and deflation techniques. We shall not pursue this issue here.

4.2. Theoretical analysis. In this subsection, we first establish a relationship between the approximations obtained from the shifted block FOM method and the shifted block GMRES method. We then show the importance of the convergence of Ritz pairs to the new algorithm. Finally, we give an error analysis on the approximate solution obtained from solving (2.5) when the inverse of A is computed inaccurately.

As the shifted block GMRES algorithm and the shifted block FOM algorithm are two popular approaches for shifted linear systems with multiple right-hand sides, we are interested in the relation between the approximations of these two methods. Denote by $\widehat{Y}_i^0$ the initial guesses, and by $\widehat{R}_i^0 = B - (I - \tau_i A^{-1})\widehat{Y}_i^0$ the residuals, $i = 1, 2, \ldots, \nu$. Let $\widehat{R}_i^0 = V_1 \widehat{R}_i$ be the QR factorization of $\widehat{R}_i^0$. Then the preconditioned and shifted block FOM algorithm uses $\widehat{Y}_i^{FOM} = \widehat{Y}_i^0 + \mathcal{V}_m \widehat{Z}_i^{FOM}$ as approximate solutions to (4.4), where

$$\widehat{Z}_i^{FOM} = (I - \tau_i \mathcal{H}_m)^{-1} E_1 \widehat{R}_i, \quad i = 1, 2, \ldots, \nu. \tag{4.10}$$

Next we consider the solutions derived from the preconditoned and shifted block GMRES algorithm [73]. Without loss of generality, let $\big(I - \tau_1 A^{-1}\big)Y_1 = B$ be the "seed" block linear system, and let $\big(I - \tau_i A^{-1}\big)Y_i = B$ $(i = 2, 3, \ldots, \nu)$ be the "additional" block linear systems [73]. In the preconditioned and shifted block GMRES algorithm, only the "seed" block system residual is minimized, whereas the residuals of the "additional" block systems are forced to be collinear with that of the "seed" block system after restarting [73]. More precisely, for the "seed" block system, the preconditioned and shifted block GMRES algorithm uses $\widehat{Y}_1^{GMRES} = \widehat{Y}_1^0 + \mathcal{V}_m \widehat{Z}_1^{GMRES}$ as an approximate solution to the "seed" system, where

$$\widehat{Z}_1^{GMRES} = \arg\min_{Z \in \mathbb{C}^{mp \times p}} \| E_1 \widehat{R}_1 - \big(\widetilde{I}_m - \tau_1 \widetilde{\mathcal{H}}_m\big) Z \|_F,$$

or equivalently,

$$\big(\widetilde{I}_m - \tau_1 \widetilde{\mathcal{H}}_m\big)^{\mathrm{H}} \big(\widetilde{I}_m - \tau_1 \widetilde{\mathcal{H}}_m\big) \widehat{Z}_1^{GMRES} = \big(\widetilde{I}_m - \tau_1 \widetilde{\mathcal{H}}_m\big)^{\mathrm{H}} E_1 \widehat{R}_1.$$

If $I - \tau_1 \mathcal{H}_m$ is nonsingular, the above equation can be reformulated as

$$\big[(I - \tau_1 \mathcal{H}_m) + |\tau_1|^2 (I - \tau_1 \mathcal{H}_m)^{-\mathrm{H}} E_m H_{m+1,m}^{\mathrm{T}} H_{m+1,m} E_m^{\mathrm{T}}\big] \widehat{Z}_1^{GMRES} = E_1 \widehat{R}_1. \tag{4.11}$$

Denote $\Gamma_1 = |\tau_1|^2 (I - \tau_1 \mathcal{H}_m)^{-\mathrm{H}} E_m$, $\Gamma_2^{\mathrm{T}} = H_{m+1,m}^{\mathrm{T}} H_{m+1,m} E_m^{\mathrm{T}}$, and

$$\Omega_1 = (I - \tau_1 \mathcal{H}_m)^{-1} \Gamma_1 \big(I + \Gamma_2^{\mathrm{T}} (I - \tau_1 \mathcal{H}_m)^{-1} \Gamma_1\big)^{-1} \Gamma_2^{\mathrm{T}}, \tag{4.12}$$

by the Sherman-Morrison-Woodbury formula [29], we have that

$$\begin{aligned} \widehat{Z}_1^{GMRES} &= (I - \tau_1 \mathcal{H}_m)^{-1} E_1 \widehat{R}_1 - \Omega_1 (I - \tau_1 \mathcal{H}_m)^{-1} E_1 \widehat{R}_1 \\ &= \widehat{Z}_1^{FOM} - \Omega_1 \widehat{Z}_1^{FOM} = (I - \Omega_1) \widehat{Z}_1^{FOM}. \end{aligned} \tag{4.13}$$

Next we consider the approximations of the "additional" block systems computed by the preconditioned and shifted block GMRES algorithm. Let $\widehat{Y}_i^{GMRES} = \widehat{Y}_i^0 + \mathcal{V}_m \widehat{Z}_i^{GMRES}$ $(i = 2, 3, \ldots, \nu)$ be the approximate solutions, and let $\widehat{R}_1^{GMRES}$ be the residual of the "seed" system. Then the residuals of the "additional" systems can be expressed as [73]

$$\widehat{R}_i^{GMRES} = \widehat{R}_1^{GMRES} W_i, \qquad i = 2, 3, \ldots, \nu, \tag{4.14}$$

where W_i are some $p \times p$ nonsingular matrices. It follows that

$$E_1\widehat{R}_i - (\widetilde{I}_m - \tau_i\widetilde{\mathcal{H}}_m)\widehat{Z}_i^{GMRES} = \big[E_1\widehat{R}_1 - (\widetilde{I}_m - \tau_1\widetilde{\mathcal{H}}_m)\widehat{Z}_1^{GMRES}\big]W_i, \quad i = 2, 3, \ldots, \nu.$$

Denote $G_1^{GMRES} = E_1\widehat{R}_1 - (\widetilde{I}_m - \tau_1\widetilde{\mathcal{H}}_m)\widehat{Z}_1^{GMRES}$, the above relation can be rewritten as

$$\left[\begin{array}{cc} \widetilde{I}_m - \tau_i\widetilde{\mathcal{H}}_m & G_1^{GMRES} \end{array}\right] \left[\begin{array}{c} \widehat{Z}_i^{GMRES} \\ W_i \end{array}\right] = E_1\widehat{R}_i, \qquad i = 2, 3, \ldots, \nu.$$

Let $\Psi_1 = G_1^{GMRES}(1 : mp, :), \Psi_2 = G_1^{GMRES}(mp+1 : mp+p, :)$ be the $mp \times p$ and $p \times p$ matrices composed of the first mp rows and the last p rows of G_1^{GMRES}, respectively. Denote by $\Phi_i = [\widetilde{I}_m - \tau_i\widetilde{\mathcal{H}}_m,\ G_1^{GMRES}]^{-1}(1 : mp, 1 : mp)$ the matrix composed of the first mp rows and mp columns of the inverse of $[\widetilde{I}_m - \tau_i\widetilde{\mathcal{H}}_m,\ G_1^{GMRES}]$. If Ψ_2 is nonsingular, then

$$\Phi_i = \big((I - \tau_i\mathcal{H}_m) + \tau_i\Psi_1\Psi_2^{-1}H_{m+1,m}E_m^{\mathrm{T}}\big)^{-1}, \quad i = 2, \ldots, \nu.$$

If $I - \tau_i\mathcal{H}_m$ is nonsingular, denote $\Psi_3^{\mathrm{T}} = \Psi_2^{-1}H_{m+1,m}E_m^{\mathrm{T}}$ and

$$\Omega_i = \tau_i(I - \tau_i\mathcal{H}_m)^{-1}\Psi_1\big(I + \tau_i\Psi_3^{\mathrm{T}}(I - \tau_i\mathcal{H}_m)^{-1}\Psi_1\big)^{-1}\Psi_3^{\mathrm{T}}, \quad i = 2, \ldots, \nu, \tag{4.15}$$

from the Sherman-Morrison-Woodbury formula [29], we obtain

$$\begin{aligned} \widehat{Z}_i^{GMRES} &= \big((I - \tau_i\mathcal{H}_m) + \tau_i\Psi_1\Psi_2^{-1}H_{m+1,m}E_m^{\mathrm{T}}\big)^{-1}E_1\widehat{R}_i \\ &= \widehat{Z}_i^{FOM} - \Omega_i\widehat{Z}_i^{FOM} = (I - \Omega_i)\widehat{Z}_i^{FOM},\ i = 2, 3, \ldots, \nu. \end{aligned} \tag{4.16}$$

In conclusion, we have the following theorem.

THEOREM 4.1. *Denote by $\widehat{Y}_i^0$ the initial guesses of the preconditioned and shifted block FOM and GMRES algorithms, and let $\widehat{Y}_i^{FOM} = \widehat{Y}_i^0 + \mathcal{V}_m\widehat{Z}_i^{FOM}$, $\widehat{Y}_i^{GMRES} = \widehat{Y}_i^0 + \mathcal{V}_m\widehat{Z}_i^{GMRES}$ $(i = 1, 2, \ldots, \nu)$ be the approximate solutions obtained from the two algorithms. Then under the above assumptions and notations, $\widehat{Z}_1^{FOM}$ and $\widehat{Z}_1^{GMRES}$ satisfy* (4.13)*, and $\widehat{Z}_i^{FOM}$ and $\widehat{Z}_i^{GMRES}$ satisfy* (4.16)*, $i = 2, \ldots, \nu$, respectively.*

Next, we show importance of the accuracy of Ritz pairs for the convergence of the preconditioned and shifted block FOM algorithm. The following theorem establishes a relation between the residuals of the shifted block linear systems and those of the Ritz pairs.

THEOREM 4.2. *Let $\widehat{R}_i^{FOM} = \tau_i V_{m+1}H_{m+1,m}(E_m^{\mathrm{T}}\widehat{Z}_i^{FOM})$, $i = 1, 2, \ldots, \nu$, be the residuals of the preconditioned and shifted block FOM algorithm. Denote by $(\mu_j, \mathcal{V}_m\mathbf{y}_j)$ the Ritz pairs of A^{-1}, and by $\mathbf{r}_j = A^{-1}\mathcal{V}_m\mathbf{y}_j - \mu_j\mathcal{V}_m\mathbf{y}_j$ $(j = 1, 2, \ldots, mp)$ the residuals of the Ritz pairs. If $\mathcal{H}_m = P\Lambda P^{-1}$ is the spectrum decomposition of $\mathcal{H}_m$, where $P = [\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_{mp}]$ and $\Lambda = diag(\mu_1, \mu_2, \ldots, \mu_{mp})$ is diagonal, then*

$$\widehat{R}_i^{FOM} = \tau_i\big[\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_{mp}\big]\big(P(I - \tau_i\Lambda)\big)^{-1}\big(E_1\widehat{R}_i\big), \quad i = 1, 2, \ldots, \nu. \tag{4.17}$$

Proof. Recall from (4.5) that the Ritz residuals can be rewritten as

$$\mathbf{r}_j = A^{-1}\mathcal{V}_m\mathbf{y}_j - \mu_j\mathcal{V}_m\mathbf{y}_j = V_{m+1}H_{m+1,m}(E_m^{\mathrm{T}}\mathbf{y}_i), \quad j = 1, 2, \ldots, mp. \tag{4.18}$$

Moreover, we have from (4.10) that

$$\widehat{R}_i^{FOM} = \tau_i V_{m+1} H_{m+1,m} E_m^{\mathrm{T}} (I - \tau_i \mathcal{H}_m)^{-1} E_1 \widehat{R}_i. \tag{4.19}$$

If $\mathcal{H}_m = P\Lambda P^{-1}$ is diagonalizable, then

$$(I - \tau_i \mathcal{H}_m)^{-1} E_1 \widehat{R}_i = P(I - \tau_i \Lambda)^{-1} P^{-1} E_1 \widehat{R}_i. \tag{4.20}$$

Thus, it follows from (4.18)–(4.20) that

$$\begin{aligned} \widehat{R}_i^{FOM} &= \tau_i V_{m+1} H_{m+1,m} E_m^{\mathrm{T}} P (I - \tau_i \Lambda)^{-1} P^{-1} E_1 \widehat{R}_i \\ &= \tau_i V_{m+1} H_{m+1,m} E_m^{\mathrm{T}} [\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_{mp}] (I - \tau_i \Lambda)^{-1} P^{-1} E_1 \widehat{R}_i \\ &= \tau_i \big[\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_{mp}\big] (I - \tau_i \Lambda)^{-1} P^{-1} E_1 \widehat{R}_i, \quad i = 1, 2, \ldots, \nu. \end{aligned}$$

□

Finally, we give an error analysis on the approximation (2.5). Let $X_i, \widetilde{X}_i$ be the "exact" and the "computed" solutions of the shifted linear systems (1.2), respectively, and denote by

$$\widetilde{Z} = \omega_0 B + \sum_{i=1}^{\nu} \widetilde{X}_i \tag{4.21}$$

the numerical approximation to $e^A B$. Then we have from (2.5) that

$$\begin{aligned} \|e^A B - \widetilde{Z}\| &= \|e^A B - \widetilde{f}^{\star}(A) + \widetilde{f}^{\star}(A) - \widetilde{Z}\| \\ &\le \|e^A B - \widetilde{f}^{\star}(A)\| + \|\widetilde{f}^{\star}(A) - \widetilde{Z}\| \\ &\le \|e^A B - \widetilde{f}^{\star}(A)\| + \sum_{i=1}^{\nu} |\omega_i| \cdot \|X_i - \widetilde{X}_i\|. \end{aligned} \tag{4.22}$$

Thus, the error is dominated by the one from the Carathéodory-Fejér approximation, and the one from the numerical solutions of the shifted linear systems. In this paper, we are interested in the latter.

Let $\widetilde{A}^{-1}$ be the "computed" solution of the exact inverse A^{-1}. Then the preconditioned shifted linear systems (4.4) turn out to be

$$(I - \tau_i \widetilde{A}^{-1}) \widetilde{Y}_i = B, \quad i = 1, 2, \ldots, \nu. \tag{4.23}$$

The following theorem shows that the quality of the approximations to the solutions of (1.2) is bounded by the relative accuracy of $\widetilde{A}^{-1}$ and the condition number of A.

THEOREM 4.3. *Let $Y_i = (I - \tau_i A^{-1})^{-1} B$, $\widetilde{Y}_i = (I - \tau_i \widetilde{A}^{-1})^{-1} B$, and $X_i = A^{-1} Y_i$, $\widetilde{X}_i = \widetilde{A}^{-1} \widetilde{Y}_i$ be the exact and the computed solutions of (1.2), respectively. Denote $F = \widetilde{A}^{-1} - A^{-1}$, if $\|\tau_i (I - \tau_i A^{-1})^{-1} F\| \ll 1$, then*

$$\frac{\|X_i - \widetilde{X}_i\|}{\|X_i\|} \lessapprox \kappa(A) \|(I - \tau_i A^{-1})^{-1}\| \cdot \frac{\|\widetilde{A}^{-1} - A^{-1}\|}{\|A^{-1}\|}, \quad i = 1, 2, \ldots, \nu. \tag{4.24}$$

Proof. We have from (4.4) and (4.23) that

$$\widetilde{Y}_i - Y_i = \tau_i (\widetilde{A}^{-1} \widetilde{Y}_i - A^{-1} Y_i) = \tau_i (\widetilde{X}_i - X_i). \tag{4.25}$$

Let $F = \widetilde{A}^{-1} - A^{-1}$, then

$$\begin{aligned}\widetilde{Y}_i - Y_i &= (I - \tau_i \widetilde{A}^{-1})^{-1} B - (I - \tau_i A^{-1})^{-1} B \\ &= \big[(I - \tau_i \big(I - \tau_i A^{-1}\big)^{-1} F\big)^{-1} - I\big](I - \tau_i A^{-1})^{-1} B \\ &= \big[(I - \tau_i \big(I - \tau_i A^{-1}\big)^{-1} F)^{-1} - I\big] Y_i.\end{aligned}$$

If $\|\tau_i (I - \tau_i A^{-1})^{-1} F\| \ll 1$, then we have that

$$\|(I - \tau_i (I - \tau_i A^{-1})^{-1} F)^{-1} - I\| = \|\tau_i (I - \tau_i A^{-1})^{-1} F\| + \mathcal{O}\big(\|\tau_i (I - \tau_i A^{-1})^{-1} F\|^2\big),$$

and

$$\frac{\|\widetilde{Y}_i - Y_i\|}{|\tau_i|} \lessapprox \|(I - \tau_i A^{-1})^{-1} F\| \cdot \|Y_i\| \le \|(I - \tau_i A^{-1})^{-1} F\| \cdot \|A\| \, \|X_i\|, \quad (4.26)$$

where we omit the high order term $\mathcal{O}\big(\|\tau_i (I - \tau_i A^{-1})^{-1} F\|^2\big)$. Combining (4.25) and (4.26), we arrive at

$$\begin{aligned}\frac{\|X_i - \widetilde{X}_i\|}{\|X_i\|} &\lessapprox \|A\| \, \|F\| \cdot \|(I - \tau_i A^{-1})^{-1}\| \\ &= \|A\| \, \|A^{-1}\| \, \|(I - \tau_i A^{-1})^{-1}\| \cdot \frac{\|F\|}{\|A^{-1}\|},\end{aligned}$$

which completes the proof. ∎

REMARK 4.2. *We notice that* $\|(I - \tau_i A^{-1})^{-1}\|_2 = \frac{1}{\sigma_{\min}(I - \tau_i A^{-1})}$, *where* $\|\cdot\|_2$ *is the 2-norm. Therefore, if* $\sigma_{\min}(I - \tau_i A^{-1})$ *is sufficiently large (say,* $\sigma_{\min}(I - \tau_i A^{-1}) \ge 10^{-4}$*), the upper bound of the relative errors of* $\{X_i\}_{i=1}^{\nu}$ *will be dominated by the condition number of* A *and the relative error of* A^{-1}*. Moreover, we point out that the assumption of* $\|\tau_i (I - \tau_i A^{-1})^{-1} F\|_2 \ll 1$ *is not stringent if* $\sigma_{\min}(I - \tau_i A^{-1})$ *is not small. Indeed, a sufficient condition is that* $\|F\|_2 \ll \frac{\sigma_{\min}(I - \tau_i A^{-1})}{|\tau_i|}$, $i = 1, 2, \ldots, \nu$.

5. Numerical Experiments. In this section, we perform some numerical experiments and show the superiority of Algorithm 2 over many state-of-the-art algorithms for computing (1.1). All the numerical experiments are run on a Dell Workstation with four core Intel(R) Pentium(R) processor with CPU 3.2 GHz and RAM 16 GB, under the Windows XP 64 bit operating system. All the experimental results are obtained from using a MATLAB 7.7 implementation with machine precision $\epsilon \approx 2.22 \times 10^{-16}$. The algorithms used in this section are listed as follows.

• **expm** is the MATLAB built-in function for matrix exponential, which implements the scaling and squaring method [34].

• **SBFOM-DR** and **PSBFOM-DR** are the shifted block FOM algorithm with deflation (Algorithm 1) and the preconditioned and shifted block FOM algorithm with deflation (Algorithm 2), respectively.

• **expv** is the MATLAB function due to Sidje [57], which evaluates $e^A \mathbf{b}$ using a restarted Krylov subspace method with a fixed dimension. The MATLAB codes are available from *http://www.maths.uq.edu.au/expokit/*.

• The MATLAB function **phipm** of Niesen and Wright [47] computes the action of linear combinations of ϕ-functions on operand vectors. The implementation combines time stepping with a procedure to adapt the Krylov subspace size. The MATLAB codes are available from *http://www1.maths.leeds.ac.uk/˜ jitse/software.html.*

• The MATLAB function **funm_kryl** is a realization of the Krylov subspace method with deflated restarting for matrix functions [20]. Its effect is to ultimately deflate a specific invariant subspace of the matrix which most impedes the convergence of the restarted Arnoldi approximation process. The MATLAB codes are available from *http://www.mathe.tu-freiberg.de/˜ guettels/funm_kryl/*.
• The MATLAB function **expmv** is developed for computing $e^A B$ [3], where A is an $n\times n$ matrix and B is $n\times p$ with $p \ll n$. It uses the scaling part of the scaling and squaring method together with a truncated Taylor series approximation to the exponential. The MATLAB codes are available from *http://www.maths.manchester.ac.uk/˜ almohy/papers.html*.

We run all these MATLAB functions with their default parameters, and the convergence tolerance in every algorithm is chosen as $tol = 10^{-8}$. In the tables below, we denote by "CPU" the CPU time in seconds, and by "Mat-Vec" the number of matrix-vector products. Let $Z(t) = e^{tA}B$ be the "exact" solution obtained from running the MATLAB build-in function `expm`, and let $\widetilde{Z}(t)$ be an approximation computed from other algorithms. Then we make use of

$$\textbf{Error} = \frac{\|Z(t) - \widetilde{Z}(t)\|_F}{\|Z(t)\|_F} \tag{5.1}$$

as the relative error of the approximation $\widetilde{Z}(t)$. If an algorithm does not converge within acceptable CPU time (say, 12 hours), or "Error" is larger than 10^{-6}, then we declare that the algorithm fails to converge. Efficiency of an algorithm is mainly measured in terms of "CPU" and "Error".

In this section, we choose $\nu = 14$ for the Carathéodory-Fejér approximation of the exponential. The right-hand sides are generated by using the MATLAB function $B =$ `randn(n,p)`, which is an n-by-p matrix with random entries. In SBFOM-DR and PSBFOM-DR, we fix the block Arnoldi steps to be $m = 30$, and set k to be 30, which is the number of approximate eigenvectors retained from the previous cycle. Except for Example 5.4, the matrix-vector product $A^{-1}\mathbf{w}$ is computed by two steps. First, we use the MATLAB command $[L, U] =$ `lu(A)` to evaluate the (sparse) LU factors of the matrix A, and then compute $A^{-1}\mathbf{w} = U \setminus (L \setminus \mathbf{w})$, where the backslash "\" is the MATLAB left matrix divide command. The CPU time of PSBFOM-DR is made up of that for computing LU decomposition, solving the shifted linear systems with multiple right-hand sides, as well as that for forming the approximation $\widetilde{Z}(t)$.

Example 5.1. This experiment is a variation of the one from Al-Mohy and Higham [3]. There are two test matrices in this example, which are generated by the MATLAB function "`gallery`". The first test matrix is the block tridiagonal matrix $A = -2500 \times$ `gallery('poisson',99)`, with $A \in \mathbb{R}^{9801\times 9801}$ being a multiple of the standard finite difference discretization of the 2D Laplacian. We want to compute $e^A B$ with the right-hand sides $B =$ `randn(9801,3)`. The second test matrix is generated by $A =$ `gallery('lesp',10000)`. It returns a $10,000 \times 10,000$ matrix with sensitive eigenvalues. We try to compute $e^A B$ with the right-hand sides $B =$ `randn(10000,5)`. Tables 5.1 and 5.2 list the numerical results.

Some remarks are in order. First, we see that PSBFOM-DR can be applied to large and sparse matrix exponential successfully. Second, the new algorithm is efficient for the computation of matrix exponential. It is seen from Tables 5.1 and 5.2 that PSBFOM-DR outperforms the other six algorithms in terms of both the CPU time and the number of matrix-vector products. Specifically, PSBFOM-DR converges much faster than SBFOM-DR, 3.17 seconds vs. 10.3 seconds for the first test problem, and 97.4 seconds vs 1318.8 seconds for the second one. This illustrates

Algorithm	CPU	Mat-Vec	Error
expm	565.0	–	–
SBFOM-DR	10.3	810	4.50×10^{-9}
PSBFOM-DR	3.17	123	5.76×10^{-10}
expv	5.98	6,944	1.39×10^{-8}
phipm	6.42	4,487	1.58×10^{-8}
funm_Kryl	9.60	900	6.83×10^{-7}
expmv	23.5	28,553	7.39×10^{-7}

Example 5.1, Table 5.1: Numerical results of the algorithms for the computation of $e^A B$, where $A = -2500 \times$ `gallery('poisson',99)` $\in \mathbb{R}^{9801 \times 9801}$, and $B =$ `randn(9801,3)`.

Algorithm	CPU	Mat-Vec	Error
expm	966.5	–	–
SBFOM-DR	1318.8	5,419	1.47×10^{-8}
PSBFOM-DR	97.4	205	3.14×10^{-9}
expv	8583.5	38,347	5.57×10^{-11}
phipm	5240.4	23,258	8.46×10^{-9}
funm_Kryl	f.c.	f.c.	f.c.
expmv	17572.3	39,225	2.73×10^{-7}

Example 5.1, Table 5.2: Numerical results of the algorithms for the computation of $e^A B$, where $A =$ `gallery('lesp',10000)` $\in \mathbb{R}^{10,000 \times 10,000}$, and $B =$ `randn(10000,5)`. Here "f.c." denotes "fails to converge".

that our preconditioning strategy is very effective for the shifted linear systems with multiple right-hand sides. Third, the number of matrix-vector products is not the whole story for computing the matrix exponential problem. For example, we notice from Table 5.1 that PSBFOM-DR used 123 matrix-vector products and 3.17 seconds, while expv used 6944 matrix-vector products and 5.98 seconds. The reason is that the CPU time of PSBFOM-DR includes that for the LU decomposition and for solving the shifted linear systems. Moreover, in the first test problem, phipm uses 4487 matrix-vector products and 6.42 seconds, while expv exploits 6944 matrix-vector products and 5.98 seconds. As a result, an algorithm using fewer matrix-vector products may not converge faster than another using more matrix-vector products, and vice versa.

Example 5.2. In this example, we try to show that the preconditioned and shifted block FOM algorithm (PSBFOM-DR) is favorable to exponential of tA with a large norm. The test matrices are two symmetric positive matrix (SPD) matrices that are available from the University of Florida Sparse Matrix Collection: *http://www.cise.ufl.edu/ research/sparse/matrices*. The first test matrix is the *1138bus* matrix arising from power system networks. It is of size 1138×1138, with 2596 nonzero elements. The second one is the *Pres_Poisson* matrix arising from computational fluid dynamics problems. The size of this matrix is $14,822 \times 14,822$, with 715,804 nonzero elements. We want to compute $e^{tA}B$, with $t = -1, -10, -100$ for the *1138bus* matrix, and $t = -100, -1000, -10000$ for the *Pres_Poisson* matrix. Tables 5.3 and 5.4 report the numerical results.

Again, the numerical results illustrate that PSBFOM-DR is superior to the state-of-the-art algorithms for the matrix exponential computation in most cases, especially when $\|tA\|$ is large. Moreover, one observes that the larger $\|tA\|$ is, the less the CPU

Algorithm	t	**CPU**	**Mat-Vec**	**Error**
expm	-1	1.31	–	–
SBFOM-DR	-1	5.28	2,190	3.17×10^{-12}
PSBFOM-DR	-1	1.95	524	4.46×10^{-12}
expv	-1	2.30	12,214	9.51×10^{-13}
phipm	-1	1.59	7,260	3.00×10^{-11}
funm_Kryl	-1	198.8	2,370	1.05×10^{-9}
expmv	-1	9.08	70,616	1.86×10^{-10}
expm	-10	1.45	–	–
SBFOM-DR	-10	180.6	74,190	7.55×10^{-12}
PSBFOM-DR	-10	0.45	164	4.07×10^{-12}
expv	-10	13.3	85,002	2.93×10^{-12}
phipm	-10	9.67	38,261	2.45×10^{-11}
funm_Kryl	-10	f.c.	f.c.	f.c.
expmv	-10	87.2	704,275	8.08×10^{-11}
expm	-100	1.59	–	–
SBFOM-DR	-100	39.2	16,050	5.38×10^{-10}
PSBFOM-DR	-100	0.47	164	1.94×10^{-11}
expv	-100	90.5	565,378	1.76×10^{-11}
phipm	-100	44.8	126,191	4.27×10^{-11}
funm_Kryl	-100	f.c.	f.c.	f.c.
expmv	-100	1109.3	7,040,894	1.47×10^{-11}

Example 5.2, Table 5.3: Numerical results of the algorithms on $e^{tA}B$, where A is the 1138×1138 *1138bus* matrix, $t = -1, -10, -100$ and $B =$ `randn(1138,4)`. Here "f.c." denotes "fails to converge".

time is required for PSBFOM-DR. The reason is that $\frac{\|(tA-\tau_i I)-tA\|}{\|tA-\tau_i I\|}$ decreases as t increases, refer to (4.1). That is, the larger $\|tA\|$ (or t) is, the better the preconditioner will be. As a comparison, the CPU time for the other algorithms increase as $\|tA\|$ (or t) becomes large. Therefore, PSBFOM-DR is preferable to the matrix exponential problem as the coefficient matrix has a large norm or when t is large.

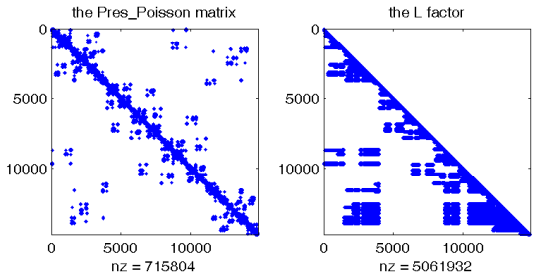

Example 5.2, Figure 5.1: Sparse structure of the *Pres_Poisson* matrix and that of the L (and U)-factor.

Algorithm	t	**CPU**	**Mat-Vec**	**Error**
expm	-100	8943.8	–	–
SBFOM-DR	-100	21.7	630	4.21×10^{-10}
PSBFOM-DR	-100	130.9	1,503	4.73×10^{-12}
expv	-100	11.2	2,139	2.16×10^{-13}
phipm	-100	7.42	1,287	7.99×10^{-12}
funm_Kryl	-100	11.9	720	1.65×10^{-12}
expmv	-100	42.2	6,411	1.60×10^{-10}
expm	-1000	8210.8	–	–
SBFOM-DR	-1000	f.c.	f.c.	f.c.
PSBFOM-DR	-1000	44.5	423	2.22×10^{-12}
expv	-1000	60.4	12,245	1.01×10^{-12}
phipm	-1000	45.7	6,746	1.80×10^{-11}
funm_Kryl	-1000	684.3	2,700	9.51×10^{-10}
expmv	-1000	410.4	61,585	1.80×10^{-10}
expm	$-10,000$	7517.6	–	–
SBFOM-DR	$-10,000$	712.3	19,050	2.96×10^{-11}
PSBFOM-DR	$-10,000$	23.7	183	1.01×10^{-13}
expv	$-10,000$	391.6	80,786	8.33×10^{-12}
phipm	$-10,000$	238.3	33,712	5.73×10^{-11}
funm_Kryl	$-10,000$	f.c.	f.c.	f.c.
expmv	$-10,000$	4095.0	613,358	4.34×10^{-11}

Example 5.2, Table 5.4: Numerical results of the algorithms on $e^{tA}B$, where A is the $14,822 \times 14,822$ *Pres_Poisson* matrix, $t = -100, -1000, -10000$ and $B =$ `randn(14822,3)`. Here "f.c." denotes "fails to converge".

Now let's briefly introduce the reason why our preconditioning strategy can be applied to large sparse matrix exponential computations successfully. In order to show this more precisely, we plot in Figure 5.1 the sparse structure of the *Pres_Poisson* matrix and that of the L-factor (which is also the U-factor) obtained from the LU factorization of A, using the MATLAB command `spy`. It is observed that the L-factor (and the U-factor) is still sparse, moreover, the expense for the LU factorization is (relatively) much lower than that for computing e^A. This explains why the cost for the LU factorization of A is not dominant in the PSBFOM-DR algorithm.

On the other hand, we notice that for the *Pres_Poisson* matrix, when $t = -1000$, PSBFOM-DR uses 44.5 seconds and 423 matrix-vector products, while phipm uses 45.7 seconds and 6746 matrix vector products to achieve the accuracy of $\mathcal{O}(10^{-12})$. That is, it seems that the number of matrix-vector products used by phipm is about 15.9 times more than PSBFOM-DR, but the CPU time of these two algorithms is about the same; see also Table 5.3 when $t = -1$. Indeed, as was mentioned in Example 5.1, the CPU time of PSBFOM-DR includes that for LU factorization of A. Furthermore, this can also be explained by using Figure 5.1, where it shows that the computational cost of performing a matrix-vector product in PSBFOM-DR is about 14 times more expensive than that in phipm.

Example 5.3. This experiment tries to illustrate that PSBFOM-DR still works well even if A is very ill-conditioned. As is known, the matrix exponential plays a fundamental role in solving linear differential equations. In this example, we consider the computation of the product of a matrix exponential with a vector, which arises

from the numerical solution of the following fractional diffusion equation [65]

$$\begin{aligned}&\frac{\partial u(x,t)}{\partial t} = d(x)\frac{\partial^\beta u(x,t)}{\partial x^\beta} + q(x,t), \quad (x,t) \in (0,1)\times(0,1], \\ &u(0,t) = 0, \quad u(1,t) = e^{-t}, \quad t \in (0,1], \\ &u(x,0) = x^3, \quad x \in (0,1),\end{aligned} \tag{5.2}$$

with the coefficient

$$d(x) = \frac{\Gamma(4-\beta)}{6}x^{1+\beta}$$

and the source term

$$q(x,t) = -(1+x)e^{-t}x^3.$$

Here $1 < \beta < 2$ and Γ is the Gamma function. For the definition of the fractional order derivative, we refer to [49].

After the spatial discretization by the shifted Grünwald formula [41], the equation (5.2) reduces to a semidiscretized ordinary differential equations of the form

$$\frac{\mathrm{d}\mathbf{u}(t)}{\mathrm{d}t} = A\mathbf{u}(t) + \mathbf{b}(t), \quad \mathbf{u}_0 = \mathbf{u}(0), \tag{5.3}$$

where $A = \frac{1}{h^\beta}DG$ with h being the grid size, D is a diagonal matrix arising from the discretization of the diffusion coefficient $d(x)$, and G is a lower Hessenberg Toeplitz matrix generated by the discretization of the fractional derivative; see [72] for the details of the discretization. The vector $\mathbf{b}(t) = e^{-t}\widetilde{\mathbf{b}}$, where $\widetilde{\mathbf{b}}$ consists of the discretization of $q(x,t)/e^{-t}$ and boundary conditions, and it is independent of t. By the variation-of-constants formula, the solution of (5.3) at time t can be expressed as

$$\begin{aligned}\mathbf{u}(t) &= e^{tA}\mathbf{u}_0 + \int_0^t e^{(t-\tau)A}e^{-\tau}\widetilde{\mathbf{b}}\mathrm{d}\tau \\ &= e^{tA}(\mathbf{u}_0 + (A+I)^{-1}\widetilde{\mathbf{b}}) - e^{-t}(A+I)^{-1}\widetilde{\mathbf{b}},\end{aligned} \tag{5.4}$$

provided that $A + I$ is invertible. In light of (5.4), to compute the solution $\mathbf{u}(t)$, we have to approximate the product of the matrix exponential e^{tA} with the vector

$$\widehat{\mathbf{b}} = \mathbf{u}_0 + (A+I)^{-1}\widetilde{\mathbf{b}},$$

which is the major computational cost for this problem.

In this experiment, we choose $\beta = 1.7$, $t = 1$, and set the size of the matrix A to be $n = 1000, 2000$ and 3000, respectively. We mention that the matrix A is very ill-conditioned and it is not a Toeplitz matrix. Indeed, as $n = 1000, 2000$ and 3000, the 1-norm condition numbers (estimated by using the MATLAB command `condest`) are about $3.44 \times 10^{10}, 3.64 \times 10^{11}$ and 1.45×10^{12}, respectively. We run the seven algorithms on this problem. Table 5.5 lists the numerical results.

We see from Table 5.5 that our new algorithm performs much better than the other algorithms in terms of CPU time and number of matrix-vector products, moreover, it can still reach an acceptable accuracy even if A is very ill-conditioned. So our new algorithm is promising even if the matrix in question is very ill-conditioned. However, it is seen that the accuracy of our approximation is lower than those obtained from

Algorithm	n	**CPU**	**Mat-Vec**	**Error**
expm	1000	9.61	–	–
SBFOM-DR	1000	f.c.	f.c.	f.c.
PSBFOM-DR	1000	1.48	41	7.11×10^{-7}
expv	1000	15.0	8,897	1.95×10^{-12}
phipm	1000	10.2	5,725	4.34×10^{-11}
funm_Kryl	1000	f.c.	f.c.	f.c.
expmv	1000	241.8	143,964	1.05×10^{-9}
expm	2000	78.3	–	–
SBFOM-DR	2000	f.c.	f.c.	f.c.
PSBFOM-DR	2000	7.66	41	1.93×10^{-7}
expv	2000	227,0	25,172	1.10×10^{-11}
phipm	2000	138.0	14,802	1.10×10^{-11}
funm_Kryl	2000	f.c.	f.c.	f.c.
expmv	2000	4196.0	472,787	1.61×10^{-9}
expm	3000	257.0	–	–
SBFOM-DR	3000	f.c.	f.c.	f.c.
PSBFOM-DR	3000	27.3	41	2.47×10^{-8}
expv	3000	1008.1	50,437	1.84×10^{-11}
phipm	3000	505.1	24,458	1.79×10^{-11}
funm_Kryl	3000	f.c.	f.c.	f.c.
expmv	3000	18749.3	946,657	2.05×10^{-9}

Example 5.3, Table 5.5, : Numerical results of Example 5.3 for computing $e^A\widehat{\mathbf{b}}$ with $\beta = 1.7$ and $n = 1000, 2000$ and 3000. The 1-norm condition numbers (estimated by using the MATLAB command `condest`) are about $3.44 \times 10^{10}, 3.64 \times 10^{11}$ and 1.45×10^{12}, respectively. Here "f.c." denotes "fails to converge".

the other algorithms. This can be interpreted by using Theorem 4.3, where it is shown that the error of the computed solution is affected by the ill-conditioning of the matrix A in question.

Example 5.4. In this example, we consider approximation of $e^{tA}B$ with A being a Toeplitz matrix. Toeplitz matrices arise from numerous topics like signal and image processing, numerical solutions of partial differential equations and integral equations, queueing networks [14, 15], and so on. The Toeplitz matrix exponential problem plays an important role in various application fields such as computational finance [36, 48, 66]. Moreover, in integral equations, the Toeplitz matrix exponential also takes part in the numerical solution of Volterra-Wiener-Hopf equations [1]. However, Toeplitz matrices generally are dense, and some classic methods for approximating the Toeplitz matrix exponential will suffer from $\mathcal{O}(n^3)$ complexities [43].

Based on the shift-and-invert Arnoldi method [44, 70], Toeplitz structure and the famous Gohberg-Semencul formula (GSF) [28], Lee, Pang and Sun [36] proposed a shift-and-invert Arnoldi algorithm for Toeplitz matrix exponential, which can reduce the computational cost to $\mathcal{O}(n \log n)$ in total. However, there are some deficiencies in this algorithm. For instance, this algorithm is a non-restarted one and there is no *posteriori* stopping criterion available, so one does not know how to restart and when to terminate this algorithm properly. Furthermore, we have no idea of choosing the *optimal* shift in the shift-and-invert Arnoldi algorithm in advance, if there is no other information *a prior*.

The aim of the example is twofold. First, we compare our preconditioned and shifted block FOM algorithm (PSBFOM-DR) with the shift-and-invert Arnoldi method for Toeplitz matrix exponential [36], and show the superiority of the former over the latter. Second, we demonstrate that PSBFOM-DR is feasible for matrix exponential of very large matrices, provided that the inverse of the matrix A can be computed efficiently. The test matrix A comes from the spatial discretization of the following fractional diffusion equation by the shifted Grünwald formula [41]

$$\frac{\partial u(x,t)}{\partial t} = d_1(x)\frac{\partial^\beta u(x,t)}{\partial_+ x^\beta} + d_2(x)\frac{\partial^\beta u(x,t)}{\partial_- x^\beta} + q(x,t),\ x \in (0,1),\ 1 < \beta < 2, \quad (5.5)$$

where $d_1(x) = 1$, $d_2(x) = 3$, and $u(0,t) = u(1,t) = 0$. We refer to [49] for the definition of the fractional order derivatives. As the diffusion coefficients are constant, the resulting matrix after the spatial discretization is a Toeplitz matrix [72].

Algorithm	n	**CPU**	**Mat-Vec**	**Error**
PSBFOM-DR	5×10^4	27.5	82	6.74×10^{-8}
SI-Arnoldi	5×10^4	28.3	120	5.09×10^{-8}
SI-BArnoldi	5×10^4	46.5	120	4.45×10^{-8}
expv	5×10^4	f.c.	f.c,	f.c.
phipm	5×10^4	f.c.	f.c,	f.c.
funm_Kryl	5×10^4	f.c.	f.c.	f.c.
expmv	5×10^4	f.c.	f.c,	f.c.
PSBFOM-DR	1×10^5	77.6	82	7.51×10^{-8}
SI-Arnoldi	1×10^5	89.2	120	1.11×10^{-7}
SI-BArnoldi	1×10^5	119.8	120	1.18×10^{-7}
expv	1×10^5	f.c.	f.c,	f.c.
phipm	1×10^5	f.c.	f.c,	f.c.
funm_Kryl	1×10^5	f.c.	f.c.	f.c.
expmv	1×10^5	f.c.	f.c,	f.c.
PSBFOM-DR	1.5×10^5	148.6	82	5.49×10^{-7}
SI-Arnoldi	1.5×10^5	169.1	120	2.40×10^{-7}
SI-BArnoldi	1.5×10^5	227.6	120	2.25×10^{-7}
expv	1.5×10^5	f.c.	f.c,	f.c.
phipm	1.5×10^5	f.c.	f.c,	f.c.
funm_Kryl	1.5×10^5	f.c.	f.c.	f.c.
expmv	1.5×10^5	f.c.	f.c,	f.c.
PSBFOM-DR	2×10^5	227.9	82	1.75×10^{-7}
SI-Arnoldi	2×10^5	261.0	120	4.10×10^{-8}
SI-BArnoldi	2×10^5	345.8	120	3.71×10^{-8}
expv	2×10^5	f.c.	f.c,	f.c.
phipm	2×10^5	f.c.	f.c,	f.c.
funm_Kryl	2×10^5	f.c.	f.c.	f.c.
expmv	2×10^5	f.c.	f.c,	f.c.

Example 5.4, Table 5.6: Numerical results of the seven algorithms for Toeplitz matrix exponential, $\beta = 1.8$, $B =$ `randn(n,2)`.

We run seven algorithms: PSBFOM-DR, expv, phipm, funm_Kryl, expmv, the shift-and-invert Arnoldi algorithm (SI-Arnoldi) [36], and a shift-and-invert block Arnoldi algorithm (SI-BArnoldi, which is a generalization of SI-Arnoldi to the block case) on

the matrices with size $n = 5 \times 10^4, 10^5, 1.5 \times 10^5$ and 2×10^5, respectively. Since A is a Toeplitz matrix, we can use the Gohberg–Semencul formula [28] for the inverse of Toeplitz matrix in PSBFOM-DR, SI-Arnoldi and SI-BArnoldi, which can be realized by using only six FFTs of length n [36, 48]. Note that one only needs to solve the linear systems only once for the Toeplitz inverse [36, 48]. For the sake of justification, we replace the matrix-vector products in PSBFOM-DR, expv, phipm, funm_Kryl, and expmv by the fast Toeplitz matrix-vector products [14, 15].

In this example, we take $\beta = 1.8$ and $t = 1$ and compute $e^{tA}B$ with B being generated by `randn(n,2)`. As was done in [36], we pick the shift $\gamma = 1/10$ in the shift-and-invert Arnoldi algorithm and the shift-and-invert block Arnoldi algorithm. As the Toplitz matrices are very large, the MATLAB build-in function `expm` is infeasible for this problem. As a compromise, we set the approximations got from running the shift-and-invert Arnoldi algorithm with the step $m = 100$ to be the "exact" solutions.

We use $m = 30$ for the PSBFOM-DR algorithm. As SI-Arnoldi is a non-restarted and non-block algorithm, we use the step $m = 60$ for SI-Arnoldi, so that the dimension of the search subspaces of PSBFOM-DR and SI-Arnoldi are the same. The number of the block shift-and-invert Arnoldi steps m is also chosen as 60 in SI-BArnoldi. Table 5.6 lists the numerical results. It is observed from this table that all the algorithms expv, phipm, funm_Kryl, and expmv fail to converge for this example. One reason is that the norm of the Toeplitz matrix A can be very large when n is large. For instance, when $n = 5 \times 10^4, 10^5, 1.5 \times 10^5$ or 2×10^5, the norm of A is in the order of $\mathcal{O}(10^9)$ or even $\mathcal{O}(10^{10})$. It is shown that the accuracy and convergence of those methods are closely related to the norm of A (refer to Example 5.2), and a large norm may lead to very slow convergence and even bad accuracy. As a comparison, PSBFOM-DR works quite well for this problem. Moreover, it converges faster than SI-Arnoldi and SI-BArnoldi in many cases, and the accuracy of the solutions obtained from Algorithm 2 is comparable to those obtained from the two shift-and-invert algorithms. These demonstrate the potential of PSBFOM-DR for computing exponential of very large Toeplitz matrices.

6. Concluding remarks. In this paper, we propose a preconditioned and shifted block FOM algorithm with deflation for computing the matrix exponential problem. The key idea for the preconditioning technique is that the absolute values of the poles of the Carathéodory-Fejér approximation can be much smaller than the norm of the matrix in question. The new method can precondition all the shifted linear systems simultaneously, and preserve the original structure of the shifted linear systems when restarting. Furthermore, when both A and B are real while the shifts are complex, the expensive step of constructing the orthogonal basis can be realized in real arithmetics.

The new algorithm can also be applied to many other problems which reduces to solving shifted linear systems with multiple right-hand sides. For instance, in the quantum chromodynamics (QCD) problems, it is common to solve linear systems with multiple shifts and multiple right-hand sides [9, 16, 69]. In the Wilson-Dirac problems, the right-hand sides represent different noise vectors, and the shifts $\{\tau_i\}_{i=1}^{\nu}$ correspond to different quark masses that are used in an extrapolation [16]. Consequently, the computation of the sign function of a large matrix A resorts to the problem of linear system which have both multiple right-hand sides and multiple complex shifts [9, 55]. In future work, inexact and restarted algorithms would be studied for solving linear systems with A in the inner iterations, where the matrix-vector products with respect to A^{-1} can be relaxed as the residual of the outer iteration approaches zero [11]. It is a very interesting topic and deserves further investigation.

Acknowledgments. The first author would like to thank Prof. Xiao-qing Jin, Prof. Hai-wei Sun and Dr. Siu-long Lei for their helpful discussions during his visiting of University of Macau.

REFERENCES

[1] M. Abdou and A. Badr, *On a method for solving an integral equation in the displacement contact problem*, Appl. Math. Comput., 127 (2002), pp. 65–78.

[2] M. Ahmad, D. Szyld, and M. van Gijzen, *Preconditioned multishift BiCG for H2-optimal model reduction*, Research Report 12-06-15, Department of Mathematics, Temple University, 2013.

[3] A. Al-Mohy and N. Higham, *Computing the action of the matrix exponential, with an application to exponential integrators*, SIAM J. Sci. Comput., 33 (2011), pp. 488–511.

[4] A. Antoulas, *Approximation of Large-Scale Dynamical Systems*, SIAM, Philadelphia, 2005.

[5] A. Archid and A. Bentbib, *Approximation of the matrix exponential operator by a structure-preserving block Arnoldi-type method*, Appl. Numer. Math., 75 (2014), pp. 37–47.

[6] M. Benzi and D. Bertaccni, *Approximate inverse preconditioning for shifted linear systems*, BIT, 43 (2003), pp. 231–244.

[7] S. Bellavia, V. Simone, D. Serafino, and B. Morini, *Efficient preconditioner updates for shifted linear systems*, SIAM J. Sci. Comput., 33 (2011), pp. 1785–1809.

[8] D. Bertaccni, *Efficient preconditioning for sequences of parameteric complex symmetric linear systems*, Electr. Tran. Numer. Anal., 18 (2004), pp. 49–64.

[9] S. Birk and A. Frommer, A CG method for multiple right-hand sides and multiple shifts in Lattice QCD calculations, arXiv preprint arXiv: 1205.0359, 2012.

[10] M. Botchev, V. Grimm, and M. Hochbruck, *Residual, restarting, and Richardson iteration for the matrix exponential*, SIAM J. Sci. Comput., 35 (2013), pp. A1376–A1397.

[11] A. Bouras and V. Frayssé, *Inexact matrix-vector products in Krylov methods for solving linear systems: a relaxation strategy*, SIAM. J. Matrix Anal. Appl., 26 (2005), pp. 660–678.

[12] A. Carpenter, A. Ruttan, and R. Varga, *Extended numerical computations on the 1/9 conjecture in rational approximation theory, in Rational Approximation and Interpolation*, Lecture Notes in Math. 1105, P. R. Graves-Morris, E. B. Saff, and R. S. Varga, eds., Springer-Verlag, Berlin, 1984, pp. 383–411.

[13] E. Celledoni and I. Moret, *A Krylov projection method for systems of ODEs*, Applied Num. Math., 24 (1997), pp. 365–378.

[14] R. Chan and M. Ng, *Conjugate gradient methods for Toeplitz systems*, SIAM Rev., 38 (1996), pp. 427–482.

[15] R. Chan and X. Jin, *An Introduction to Iterative Toeplitz Solvers*, SIAM, Philadelphia, 2007.

[16] D. Darnnu, R. Morgan, and W. Wilcox, *Deflated GMRES for systems with multiple shifts and multiple right-hand sides*, Lin. Alg. Appl., 429 (2008), pp. 2415–2434.

[17] V. Druskin and L. Knizhnerman, *Krylov subspace approximation of eigenpairs and matrix functions in exact and computer arithmetic*, Numer. Linear Algebra Appl., 2 (1995), pp. 205–217.

[18] V. Druskin and L. Knizhnerman, *Extended Krylov subspaces: Approximation of the matrix square root and related functions*, SIAM J. Matrix Anal. Appl., 19 (1998), pp. 755–771.

[19] M. Eiermann and O. Ernst, *A restarted Krylov subspace method for the evaluation of matrix functions*, SIAM J. Numer. Anal., 44 (2006), pp. 2481–2504.

[20] M. Eiermann, O. Ernst, and S. Güttel, *Deflated restarting for matrix functions*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 621–641.

[21] A. Frommer, *BiCGStab(ν) for families of shifted linear systems*, Computing, 70 (2003), pp. 87–109.

[22] A. Frommer and U. Glassner, *Restarted GMRES for shifted linear systems*, SIAM J. Sci. Comput., 19 (1998), pp. 15–26.

[23] A. Frommer, K. Kahl, T. Lippert and H. Rittich, *2-norm error bounds and estimates for Lanczos approximations to linear systems and rational matrix functions*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 1046–1065.

[24] A. Frommer and V. Simoncini, *Stopping criteria for rational matrix functions of Hermitian and symmetric matrices*, SIAM J. Sci. Comput, 30 (2008), pp. 1387–1412.

[25] R. Freund, *On conjugate gradient type methods and polynomial preconditioners for a class of complex non-Hermitian matrices*, Numer. Math., 57 (1990), pp. 285–312.

[26] E. Gallopoulos and Y. Saad, *Efficient solution of parabolic equations by polynomial approximation methods*, SIAM J. Stat. Sci. Comput., 13 (1992), pp. 1236–1264.

[27] R. Garrappa and M. Popolizio, *On the use of matrix functions for fractional partial differential equations*, Mathematics and Computers in Simulation, 81 (2011), pp. 1045–1056.

[28] I. Gohberg and A. Semencul, *On the inversion of finite Toeplitz matrices and their continuous analogs*, Mat. Issled., 2 (1972), pp. 201–233.

[29] G.H. Golub and C.F. Van Loan, *Matrix Computations*, 4th ed., John Hopkins University Press, Baltimore, MD, 2012.

[30] G. Gu, X. Zhou, and L. Lin, *A flexible preconditioned Arnoldi method for shifted linear systems*, J. Comput. Math., 25 (2007), pp. 522–530.

[31] B. Jegerlehner, *Krylov space solvers for shifted linear systems*, Arxiv preprint hep-lat/9612014, 1996.

[32] W. Jiang and Gang Wu, *A thick-restarted block Arnoldi algorithm with modified Ritz vectors for large eigenproblems*, Comput. Math. Appl., 60 (2010), pp. 873–889.

[33] N.J. Higham, *Functions of Matrices: Theory and Computation*, SIAM, Philadelphia, 2008.

[34] N.J. Higham, *The scaling and squaring method for the matrix exponential revisited*, SIAM Rev., 51 (2009), pp. 747–764.

[35] M. Hochbruck and A. Ostermann, *Exponential integrators*, Acta Numer., 19 (2010), pp. 209–286.

[36] S. Lee, H. Pang, and H. Sun, *Shift-invert Arnoldi approximation to the Toeplitz matrix exponential*, SIAM J. Sci. Comput., 32 (2010), pp. 774–792.

[37] L. Lopez and V. Simoncini, *Analysis of projection methods for rational function approximation to the matrix exponential*, SIAM J. Numer. Anal., 44 (2006), pp. 613–635.

[38] L. Lopez and V. Simoncini, *Preserving geometric properties of the exponential matrix by block Krylov subspace methods*, BIT, 46 (2006), pp. 813–830.

[39] A. Magnus, *Asymptotics and super asymptotics for best rational approximation error norms to the exponential function (the "1/9" problem) by the Carathéodory-Fejér method*, Nonlinear Numer. Methods Ration. Approx. II 296 (1994), pp. 173–185.

[40] MATLAB 7, The MATHWORKS INC., 2004.

[41] M. Meerschaert and C. Tadjeran, *Finite difference approximations for fractional advection-dispersion flow equations*, J. Comput. Appl. Math., 172 (2004), pp. 65–77.

[42] G. Meinardus, *Approximation of Functions: Theory and Numerical Methods*, Springer-Verlag New York, Inc., New York, 1967.

[43] C. Moler and C.F. Van Loan, *Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later*, SIAM Rev., 45 (2003), pp. 3–49.

[44] I. Moret and P. Novati, *RD-rational approximations of the matrix exponential*, BIT, 44 (2004), pp. 595–615.

[45] I. Moret and M. Popolizio, *The restarted shift-and-invert Krylov method for matrix functions*, Numer. Linear Algebra Appl., 21 (2014), pp. 68–80.

[46] R. Morgan, *Restarted block GMRES with deflation of eigenvalues*, Appl. Numer. Math., 54 (2005), pp. 222–236.

[47] J. Niesen and W. Wright, *Algorithm 919: A Krylov subspace algorithm for evaluating the ϕ-functions appearing in exponential integrators*, ACM Trans. Math. Software, 38 (3), Article 22, 2012.

[48] H. Pang and H. Sun, *Shift-Invert Lanczos method for the symmetric positive semidefinite Toeplitz matrix exponential*, Numer. Linear Algebra Appl., 18 (2011), pp. 603–614.

[49] I. Podlubny, *Fractional Differential Equations*, Academic Press, New York, 1999.

[50] M. Popolizio and V. Simoncini, *Acceleration techniques for approximating the matrix exponential operator*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 657–683.

[51] A. Quarteroni and A. Valli, *Numerical Approximation of Partial Differential Equations*. Springer, Berlin, 1997.

[52] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd edition, SIAM, Philadelphia, 2003.

[53] Y. Saad, *Analysis of some Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal., 29 (1992), pp. 209–228.

[54] A. Saibaba, T. Bakhos and P. Kitanidis, *A flexible Krylov solver for shifted systems with application to oscillatory hydraulic tomography*, SIAM J. Sci. Comput., 35 (2013), pp. A3001–A3023.

[55] T. Sakurai, H. Tadano and Y. Kuramashi, *Application of block Krylov subspace algorithms to the Wilson–Dirac equation with multiple right-hand sides in lattice QCD*, Comput. Phy. Commun., 181 (2010), pp. 113–117.

[56] T. Schmelzer and L. N. Trefethe, *Evaluating matrix functions for exponential integrators via Carathéodory-Fejér approximation and contour integrals*, Electr. Tran. Numer. Anal., 29 (2007), pp. 1–18.

[57] R. Sidje, *EXPOKIT: Software Package for Computing Matrix Exponentials*, ACM Tran. Math. Soft., 24 (1998), pp. 130–156.

[58] V. Simoncini, *Restarted full orthogonalization method for shifted linear systems*, BIT, 43 (2003), pp. 459–466.

[59] V. Simoncini, *The extended Krylov subspace for parameter dependent systems*, Applied Num. Math., 60 (2010), pp. 550–560.

[60] V. Simoncini and D. Szyld, *Recent computational developments in Krylov Subspace methods for linear systems*, Numerical Linear Algebra, 14 (2007), pp. 1–59.

[61] V. Simoncini and D. Szyld, *Theory of inexact Krylov subspace methods and applications to scientific computing*, SIAM J. Sci. Comput., 25 (2003), pp. 454–477.

[62] G. Sleijpen, J. van den Eshof and M. van Gijzen, *Restarted GMRES with inexact matrix-vector products*, Technical Report TR/PA/04/75, CERFACS, Toulouse, France, 2004.

[63] K. Soodhalter, D. Szyld, and F. Xue, *Krylov subspace recycling for sequences of shifted linear systems*, Research Report 13-01-10, Department of Mathematics, Temple University, 2013.

[64] A. Suárez, H. Sarmiento, E. Flórez, M.D. García, and G. Montero, *Updating incomplete factorization preconditioners for shifted linear systems arising in a wind model*, J. Comput. Appl. Math., 225 (2011), pp. 2640–2646.

[65] C. Tadjeran, M. Meerschaert, and H. Scheffler, *A second-order accurate numerical approximation for the fractional diffusion equation*, J. Comput. Phy., 213 (2006), pp. 205–213.

[66] D. Tangman, A. Gopaul and M. Bhuruth, *Exponential time integration and Chebychev discretisation schemes for fast pricing of options*, Appl. Numer. Math., 58 (2008), pp. 1309–1319.

[67] L. Trefethen and M. Gutknecht, *The Carathéodory-Fejér method for real rational approximation*, SIAM J. Numer. Anal., 20 (1983), pp. 420–436.

[68] L. Trefethen, J. Weideman and T. Schmelzer, *Talbot quadratures and rational approximations*, BIT, 46 (2006), pp. 653–670.

[69] J. van den Eshof, A. Frommer, T. Lippert, K. Schilling, and H.A. van der Vorst, *Numerical methods for the QCD overlap operator. I. Sign-function and error bounds*, Comput. Phy. Commun., 146 (2002), pp. 203–224.

[70] J. van den Eshof and M. Hochbruck, *Preconditioning Lanczos approximations to the matrix exponential*, SIAM J. Sci. Comput., 27 (2006), pp. 1438–1457.

[71] J. van Deun and L. N. Trefethen, *A robust implementation of the Carathédory-Fejér method for rational approximation*, BIT, 51: 1039–1050, 2011.

[72] H. Wang, K. Wang, and T. Sircar, *A direct $O(N \log^2 N)$ finite difference method for fractional diffusion equations*, J. Comput. Phys., 229 (2010), pp. 8095–8104.

[73] G. Wu, Y. Wang, and X. Jin, *A preconditioned and shifted GMRES algorithm for the PageRank problem with multiple damping factors*, SIAM J. Sci. Comput., 34 (2012), pp. A2558–A2575.

[74] G. Wu and Y. Wei, *On analysis of projection methods for rational function approximation to the matrix exponential*, SIAM J. Numer. Anal., 48 (2010), pp. 191–197.

[75] Q. Ye, *Error bounds for the Lanczos methods for approximating matrix exponentials*, SIAM J. Numer. Anal., 51 (2013), pp. 68–87.